

(12) **United States Patent**
Humby et al.

(10) **Patent No.:** **US 9,063,940 B1**
(45) **Date of Patent:** **Jun. 23, 2015**

(54) **SUPERSEDING OBJECTS IN A RETENTION SYSTEM**

(75) Inventors: **David Humby**, Richmond (CA); **Fiona Schrader**, Ottawa (CA); **Steve Sauder**, Ottawa (CA); **David Mennie**, Ottawa (CA); **John-David Dorman**, Ottawa (CA)

(73) Assignee: **EMC Corporation**, Hopkinton, MA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 759 days.

(21) Appl. No.: **11/541,485**

(22) Filed: **Sep. 29, 2006**

(51) **Int. Cl.**
G06F 17/30 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 17/30085** (2013.01); **G06F 17/3023** (2013.01)

(58) **Field of Classification Search**
CPC G06F 17/30085; G06F 17/3023
USPC 707/694
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,806,078 A * 9/1998 Hug et al. 715/205
5,864,875 A * 1/1999 Van Huben et al. 1/1

6,950,257 B1 *	9/2005	Greco et al.	360/48
7,404,140 B2 *	7/2008	O'Rourke	715/222
2004/0153965 A1 *	8/2004	O'Rourke	715/505
2004/0167901 A1 *	8/2004	Margolus	707/10
2005/0055518 A1 *	3/2005	Hochberg et al.	711/159
2005/0055519 A1 *	3/2005	Stuart et al.	711/159
2005/0071390 A1 *	3/2005	Midgley et al.	707/204
2005/0273858 A1 *	12/2005	Zadok et al.	726/24
2006/0010150 A1 *	1/2006	Shaath et al.	707/102
2007/0094311 A1 *	4/2007	Pelletier et al.	707/204
2007/0130232 A1 *	6/2007	Therrien et al.	707/204

* cited by examiner

Primary Examiner — James Trujillo

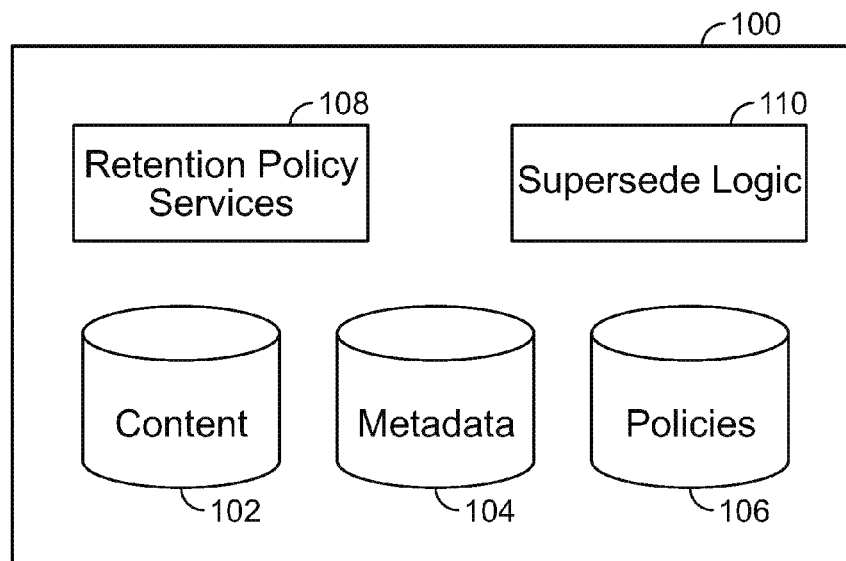
Assistant Examiner — Fatima Mina

(74) *Attorney, Agent, or Firm* — Van Pelt, Yi & James LLP

(57) **ABSTRACT**

Superseding a prior version of a document, to which prior version a retention policy or other requirement has been applied, is disclosed. In some embodiments, an attribute of a retention policy indicates whether a document to which the retention policy has been applied is to be superseded by a subsequently created and/or saved version of the document. In some embodiments, the attribute is set by a logic or process configured to apply the retention policy to the document. If the retention policy indicates that supersede is enabled, in various embodiments when a subsequent version is created and/or saved, the prior version is promoted to the final phase of the retention policy that has been applied to it and automatically “qualified” for disposition as indicated in the final phase of the retention policy, without regard to intervening requirements, processes, phases, approvals, retention, waiting, or other periods, etc.

21 Claims, 5 Drawing Sheets



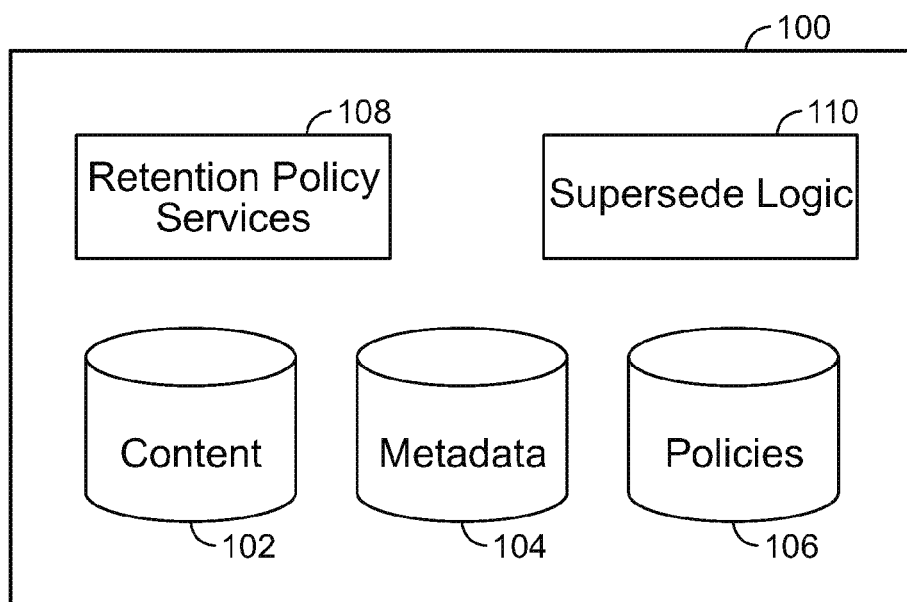


FIG. 1

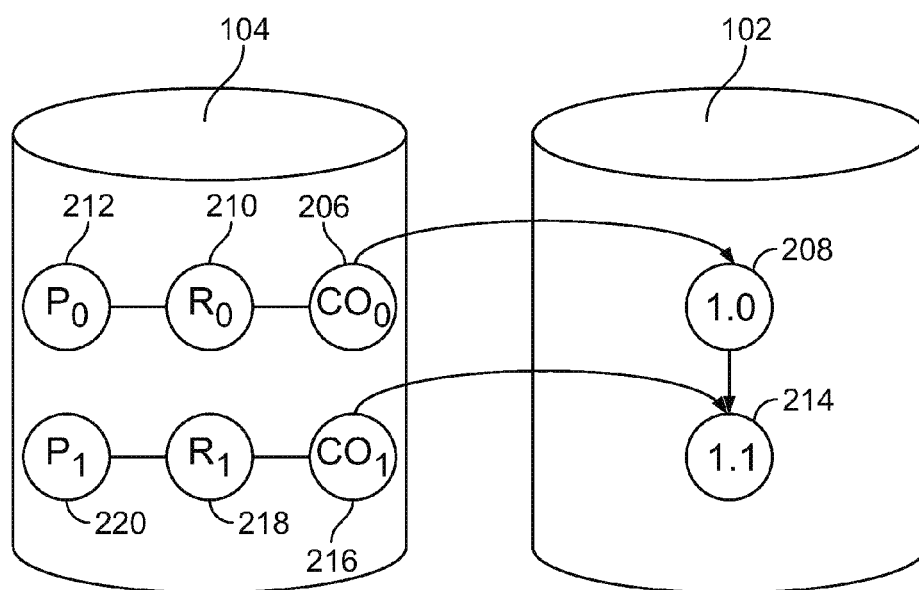


FIG. 2

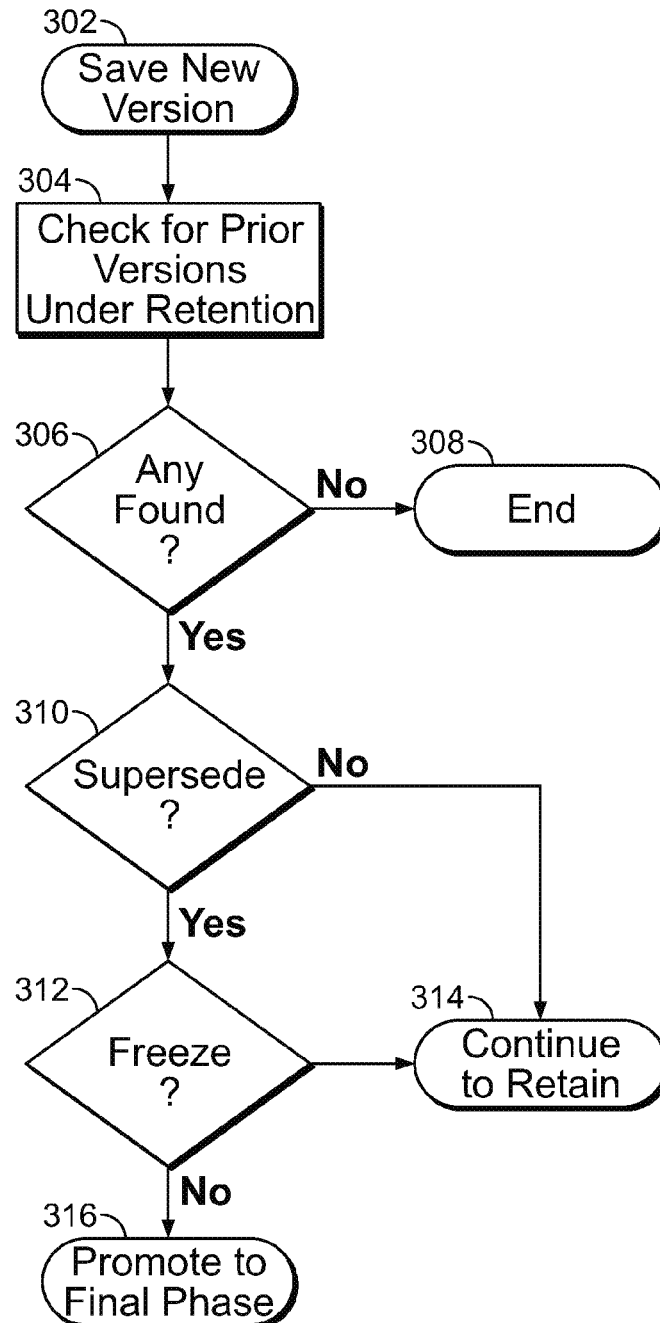


FIG. 3

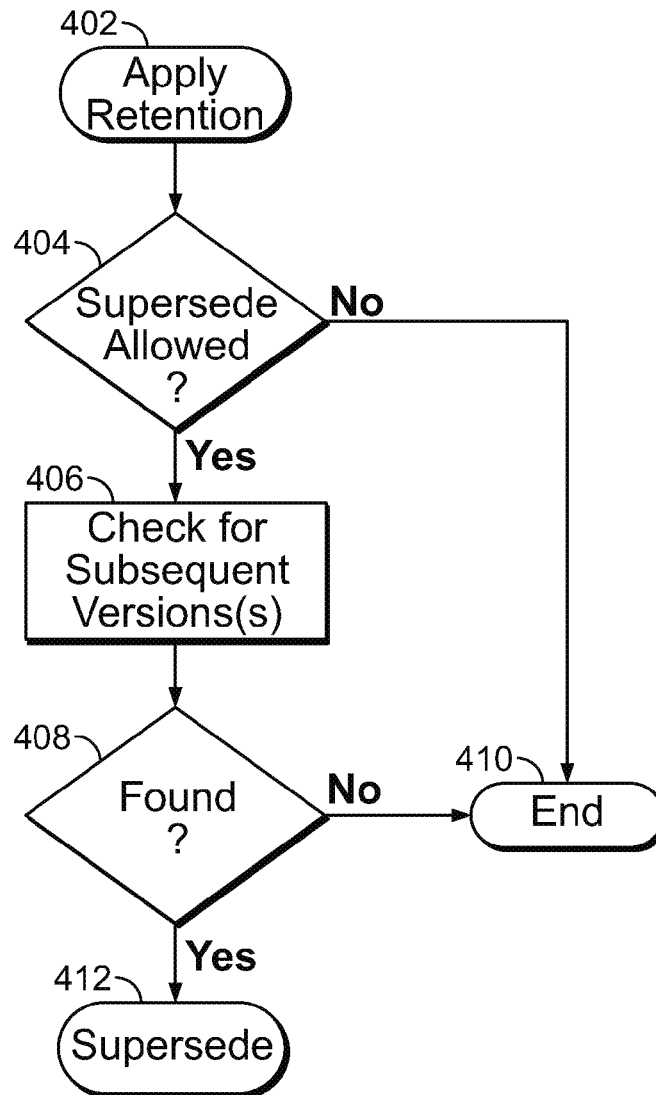


FIG. 4

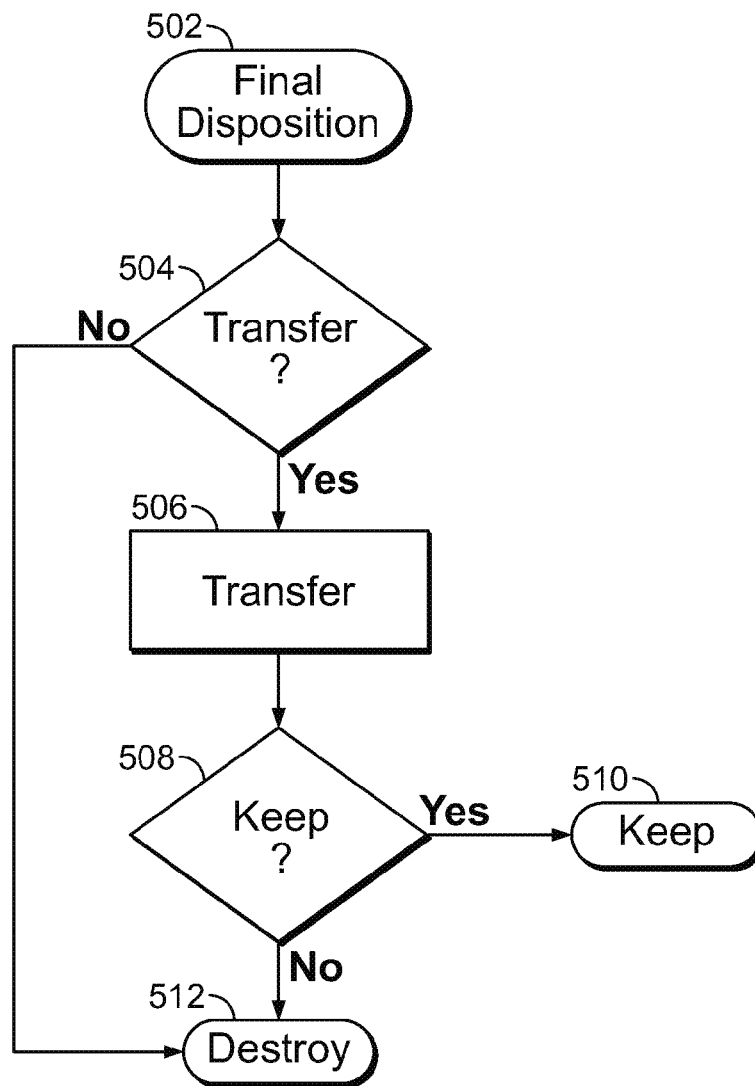


FIG. 5

SUPERSEDING OBJECTS IN A RETENTION SYSTEM

BACKGROUND OF THE INVENTION

Records management system, content management system, and other solutions have been provided to manage and provide access to stored content. One aspect managed by some such systems is document/record retention, including disposition (e.g., destruction) once all applicable retention periods and/or requirements have ended. However, for some content objects, it may not always (or in some cases ever) be required and/or desirable to retain all prior versions of a document or other stored content object. For example, prior versions created during a particular phase of a business process to which the document relates, such as an initial creation phase during which an originating author of the document might modify the document many times, with minor changes in each successive version, may not be required or desired to be retained. In such cases, there may be a need to ensure the most recent version is retained, and one way of doing so would be to apply a retention policy to the document, with the retention being inherited by subsequent version. Typically, however, applying retention to a document results in that document being retained, even after a subsequent version has been created and saved. In some cases, there may be a need to be able to discard older versions of less critical content, e.g., to reclaim associated storage space to be used to store more recent and/or more important content. Therefore, there is a need to be able to override retention of at least some prior versions of content objects to which retention has been applied.

BRIEF DESCRIPTION OF THE DRAWINGS

Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings.

FIG. 1 is a block diagram illustrating an embodiment of a content management system.

FIG. 2 is a block diagram illustrating an embodiment of retention as applied to successive versions of a content item stored in a content management system.

FIG. 3 is a flow chart illustrating an embodiment of a process for selectively superseding prior versions of content items to which retention has been applied.

FIG. 4 is a flow chart illustrating an embodiment of a process for selectively superseding an existing version of a content item to which existing version a retention policy with respect to which supersede is enabled is applied.

FIG. 5 is a flow chart illustrating an embodiment of a process for final disposition in accordance with an applicable retention policy.

DETAILED DESCRIPTION

The invention can be implemented in numerous ways, including as a process, an apparatus, a system, a composition of matter, a computer readable medium such as a computer readable storage medium or a computer network wherein program instructions are sent over optical or communication links. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. A component such as a processor or a memory described as being configured to perform a task includes both a general component that is temporarily configured to perform the task at a given time or a specific component that is

manufactured to perform the task. In general, the order of the steps of disclosed processes may be altered within the scope of the invention.

A detailed description of one or more embodiments of the invention is provided below along with accompanying figures that illustrate the principles of the invention. The invention is described in connection with such embodiments, but the invention is not limited to any embodiment. The scope of the invention is limited only by the claims and the invention encompasses numerous alternatives, modifications and equivalents. Numerous specific details are set forth in the following description in order to provide a thorough understanding of the invention. These details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

Superseding a prior version of a document, to which prior version a retention policy or other requirement has been applied, is disclosed. In some embodiments, an attribute of a retention policy instance and/or associated object indicates whether a document or other content item to which the retention policy has been applied is to be superseded by a subsequently created and/or saved version of the document. In some embodiments, the attribute is set by a retention or other logic or process configured to apply the retention policy to the document. If the retention policy indicates that supersede is enabled, in various embodiments when a subsequent version is created and/or saved, the prior version is promoted to the final phase of the retention policy that has been applied to it and automatically “qualified” for disposition as indicated in the final phase of the retention policy, without regard to intervening requirements, processes, phases, approvals, retention, waiting, or other periods, etc. Final disposition is determined and performed in accordance with the retention policy.

FIG. 1 is a block diagram illustrating an embodiment of a content management system. In the example shown, the content management system 100 includes a content store 102 configured to store one or more documents and/or other content items, a metadata store 104 configured to store metadata associated with the content items stored in content store 102. In some embodiments, each content item in content store 102 is represented by one or more associated objects stored in metadata store 104. In some embodiments, a content server, not shown in FIG. 1, provides access to content items stored in content store 102 and/or metadata in metadata store 104. In some embodiments, metadata in metadata store 104 may be used to located one or more content items stored in content store 102, e.g., based on keywords or other attributes (e.g., author, date created, etc.) of the content items. One or more retention policies are stored in a policy store 106. Retention policy services 108 are invoked by application and/or content management system business logic, as appropriate, to enforce policies that have been applied to corresponding content items in content store 102. In some embodiments, a retention (or other) policy is applied to a content item by associating a retainer object and/or an instance of the retention policy with a metadata object that represents the content item in metadata stored in metadata store 104, as described further below with reference to FIG. 2. The content management system 100 also includes a supersede logic 110. Supersede logic 110 in some embodiments overrides a retention policy in prescribed circumstances, e.g., by automatically promoting to a final phase of retention, and qualifying for final disposition at that final phase, a document or other content item that has been super-

seded by a subsequent version, e.g., when supersede is enabled with respect to the prior version.

FIG. 2 is a block diagram illustrating an embodiment of retention as applied to successive versions of a content item stored in a content management system. In the example shown, a first version (1.0) of a content item **208** stored in content store **102** is represented in metadata store **104** by a content object **206**. Retention has been applied to the content item **208** by associating with the content object **206** a retainer object **210** configured to apply with respect to the content object **206** and its associated content item **208** a retention policy instance **212**. In the example shown, a second, subsequent version (1.1) of the content item, i.e., content item **214** has been created and saved in content store **102**. The new version **214** is represented in metadata by an associated content object **216**. Retention has been applied by associating with the content object **216** a retainer object **218** configured to apply with respect to the content object **216** and its associated content item **214** a retention policy instance **220**. In some embodiments, if the retention policy instance **212** or some other source and/or information indicated that a subsequent version, if any, should supersede the content item **208**, i.e., supersede is enabled with respect to the content item **208** and/or retention policy instance **212**, the content item **208** would be superseded upon the creation and/or saving (depending on the embodiment) of new version **214**. In some embodiments, supersede logic **110** would be invoked and would promote content item **208** to the final phase of retention under policy instance **212**, and would immediately qualify content item **208** for final disposition. In some embodiments, content item **208** would be qualified for final disposition without regarding to any intervening phases, retention or other time period, review and/or approval requirements, etc. In various embodiments, final disposition may include one or more of the following: transfer to offline storage and deletion from the content management system, transfer and retention within the content management system, and total destruction (i.e., removal from the content management system without transfer).

FIG. 3 is a flow chart illustrating an embodiment of a process for selectively superseding prior versions of content items to which retention has been applied. In some embodiments, the process of FIG. 3 is implemented by supersede logic **110** of FIG. 1. In the example shown, when a new version of a content item is saved (**302**), a check is done to determine whether any prior version of the content item is under retention (**304**). If no prior versions under retention are found (**306**), the process ends (**308**). If one or more prior versions under retention are found (**306**), it is determined with respect to each such prior version whether supersede is enabled with respect to that prior version (**310**). In some embodiments, an attribute of an associated retention policy instance is checked to determine whether supersede is enable. If supersede is not enabled (**310**), retention of the prior version continues as before (**314**). If supersede is enabled, a check is done to determine if a freeze has been placed on the content item (**312**), e.g., due to litigation, regulatory action, applicability to an ongoing investigation, etc. If a freeze is in place (**312**), retention of the prior version continues as before (**314**). If supersede is enabled and no hold is in place (**310**, **312**), the content item is promoted to and qualified for final disposition in a final phase of the retention policy that has been applied to it (**316**). During disposition processing, if a hold is applied to the content object, even though the item has been superseded, the content object cannot be destroyed.

FIG. 4 is a flow chart illustrating an embodiment of a process for selectively superseding an existing version of a

content item to which existing version a retention policy with respect to which supersede is enabled is applied. In some embodiments, the process of FIG. 4 is implemented by supersede logic **110** of FIG. 1. In the example shown, when a retention policy is applied to a content item (**402**), a check is performed to determine whether supersede is enabled (**404**). If not, the process ends (**410**). If supersede is enabled (**404**), a check for subsequent versions of the content item is performed (**406**). If no subsequent version is found (**408**), the process ends. If a subsequent version is found (**408**), the content item to which the retention policy was applied (**402**) is superseded (**412**). As described above, the superseded content item is promoted to the final phase of the applied retention policy, regardless of the conditions, requirements, and intervening phases of the retention policy, and immediately qualified for final disposition as provided for in the policy. The process of FIG. 4 would allow, for example, for a retention policy to be applied to a whole type and/or class of content item (e.g., those stored in a particular physical and/or logical storage location, such as a folder or subdirectory) without resulting in retention of older versions that are no longer required and/or desired to be retained.

FIG. 5 is a flow chart illustrating an embodiment of a process for final disposition in accordance with an applicable retention policy. An indication is received that a content item has qualified for final disposition in accordance with an applicable retention policy (**502**). If the final disposition provides for transfer to offline or other storage (**504**), the transfer is performed (**506**). After transfer (**506**), the content item is either retained in the content management system (**510**) or removed (**512**), depending on the policy. If the retention policy calls for "total destruction" at final disposition, the content item is removed from the content management system without first transferring a copy to another storage system, location, and/or media volume (**504**, **512**). In some embodiments, whatever final disposition the retention policy indicates, an audit trail is generated to document the final disposition. In some embodiments, the audit trail includes, if applicable, information indicating that the content item qualified for final disposition as a result of being superseded as described herein.

Superseding unwanted and/or unneeded prior versions of content items has been described. Superseding older version that are not required and/or desired to be retained conserves storage space and ensures that only information that is currently relevant and/or required to be retained is stored and managed.

Although the foregoing embodiments have been described in some detail for purposes of clarity of understanding, the invention is not limited to the details provided. There are many alternative ways of implementing the invention. The disclosed embodiments are illustrative and not restrictive.

What is claimed is:

1. A method of managing stored content, comprising:

using one or more processors to apply a retention policy to a content item, the retention policy comprising one or more phases for the content item to go through including a final phase that indicates a final disposition for the content item, and wherein applying the retention policy to the content item includes associating the content item with a retention policy instance that includes an attribute that indicates whether a supersede is enabled with respect to the content item;

determining that a subsequent version of the content item exists;

using the one or more processors to determine whether the attribute associated with the applied retention policy

5

instance indicates that the supersede with respect to the content item is enabled, wherein the supersede allows the content item to be superseded by the subsequent version; and

in the event that the attribute indicates that supersede with respect to the content item is enabled, using the one or more processors to:

override the retention policy applied to the content item so as to promote the content item to the final phase of the retention policy instance associated with the content item and immediately qualify the content item for final disposition, without regard to any intervening requirement or phase of the retention policy instance associated with the content item.

2. A method as recited in claim 1, wherein determining that the supersede with respect to the retention policy is enabled includes reading an attribute or other data comprising or otherwise associated with the retention policy.

3. A method as recited in claim 2, wherein reading an attribute or other data comprising or otherwise associated with the retention policy includes reading an attribute of the instance of the retention policy associated with the content item.

4. A method as recited in claim 3, wherein the attribute of the instance of the retention policy associated with the content item is set by one or more of the following: the retention policy, business logic, and application logic.

5. A method as recited in claim 1, wherein the promoting the content item to the final phase of the retention policy includes promoting the content item to the final phase as indicated in the instance of a retention policy associated with the content item.

6. A method as recited in claim 1, wherein final disposition includes generating an audit trail.

7. A method as recited in claim 6, wherein the audit trail reflects that the content item was qualified for final disposition as a result of being superseded.

8. A method as recited in claim 1, wherein the retention policy is applied or determined to be required to be applied to the content item at a time subsequent to creation of the subsequent version.

9. A method as recited in claim 1, wherein the final disposition includes one or more of the following: transfer to a storage location or media outside a content management system in which the content item has been stored; removal of the content item from the content management system subsequent to transfer; and total destruction of the content item through removal of the content item from the content management system without prior transfer.

10. A content management system, comprising:

a storage area configured to store a content item to which a retention policy has been or is to be applied, the retention policy comprising one or more phases for the content item to go through including a final phase that indicates a final disposition for the content item, and wherein applying the retention policy to the content item includes associating the content item with a retention policy instance that includes an attribute that indicates whether a supersede is enabled with respect to the content item; and

a hardware processor coupled to the storage area, the processor configured to:

determine that a subsequent version of the content item exists;

determine whether the attribute associated with the applied retention policy instance indicates that the supersede with respect to the content item is enabled,

6

wherein the supersede allows the content item to be superseded by the subsequent version; and

in the event that the attribute indicates that supersede with respect to the content item is enabled, override the retention policy applied to the content item so as to promote the content item to the final phase of the retention policy instance associated with the content item and immediately qualify the content item for final disposition, without regard to any intervening requirement or phase of the retention policy instance associated with the content item.

11. A content management system as in claim 10, wherein the processor is configured to determine that the supersede with respect to the retention policy is enabled includes at least in part by reading an attribute or other data comprising or otherwise associated with the retention policy.

12. A content management system as in claim 11, wherein reading an attribute or other data comprising or otherwise associated with the retention policy includes reading an attribute of the instance of the retention policy associated with the content item.

13. A content management system as in claim 12, wherein the attribute of the instance of the retention policy associated with the content item is set by one or more of the following: the retention policy, business logic, and application logic.

14. A content management system as in claim 10, wherein final disposition includes generating an audit trail that reflects that the content item was qualified for final disposition as a result of being superseded.

15. A computer program product for managing stored content, the computer program product being stored on a non-transitory computer readable medium and comprising computer instructions for:

applying a retention policy to the content item, the retention policy comprising one or more phases for the content item to go through including a final phase that indicates a final disposition for the content item, and wherein applying the retention policy to the content item includes associating the content item with a retention policy instance that includes an attribute that indicates whether a supersede is enabled with respect to the content item;

determining that a subsequent version of the content item exists;

determining whether the attribute associated with the applied retention policy instance indicates that the supersede with respect to the content item is enabled, wherein the supersede allows the content item to be superseded by the subsequent version; and

in the event that the attribute indicates that supersede with respect to the content item is enabled, overriding the retention policy applied to the content item so as to promote the content item to the final phase of the retention policy instance associated with the content item and immediately qualifying the content item for final disposition, without regard to any intervening requirement or phase of the retention policy instance associated with the content item.

16. A computer program product as recited in claim 15, wherein determining that the supersede with respect to the retention policy is enabled includes reading an attribute or other data comprising or otherwise associated with the retention policy.

17. A computer program product as recited in claim 15, wherein final disposition includes generating an audit trail that reflects that the content item was qualified for final disposition as a result of being superseded.

18. A method as recited in claim 1, further comprising determining if a freeze has been placed on the content item.

19. The method as recited in claim 1, wherein in the event that the attribute indicates that supersede with respect to the content item is not enabled, continue to retain the content item 5 according to the applied retention policy.

20. The system as recited in claim 10, wherein in the event that the attribute indicates that supersede with respect to the content item is not enabled, continue to retain the content item according to the applied retention policy. 10

21. The computer program product as recited in claim 15, wherein in the event that the attribute indicates that supersede with respect to the content item is not enabled, continue to retain the content item according to the applied retention policy. 15

* * * * *